

A Submodular Approach to Acoustic Sensor Placement using the MARLIN Digital Twin

Edward Clark¹, Eric Schoof², Chris Manzie², and Alan Hunter¹

¹University of Bath, Bath BA2 7AY, United Kingdom

²University of Melbourne, Victoria 3010, Australia

Contact author: Edward Clark, eprc20@bath.ac.uk

Abstract: We present a submodular optimization approach for acoustic sensor placement using the MARLIN digital twin. Our objective is to select sensor locations in a region off the coast of Victoria, Australia, to maximize detection performance across the environment. The sensor placement problem is formulated as a submodular set function, allowing the use of a greedy algorithm that guarantees solutions within a known bound of the optimum. MARLIN provides a high-fidelity simulation environment using open-source oceanographic and bathymetric data. We generate 2,840 candidate sensor locations and compute acoustic propagation loss using the PyRAM model, incorporating environmental variability. Detection probability is estimated from signal-to-noise ratios, and the joint detectability across the region is evaluated for different sensor configurations. We compare sequential and simultaneous sensor placement strategies, analysing their impact on detection performance and computational efficiency. Results show that the greedy submodular approach achieves near-optimal detection with orders-of-magnitude fewer evaluations than brute-force search, making it practical for large-scale deployments. The study demonstrates the effectiveness of submodular optimization for sensor placement in complex marine environments and discusses trade-offs between placement strategies.

Keywords: *Passive acoustics, Sensor Placement, Submodular Selection*

1. INTRODUCTION

Passive acoustic underwater sensors are key to providing monitoring and deterrence in ocean environments. Their ability for long-range sensing makes them suitable for a range of applications; however, complex acoustic propagation makes optimal placement of them challenging. Temperature, salinity, and pressure gradients all influence the way sound travels through the water column, leading to sound channels where sound can travel hundreds of kilometres, and ‘shadow’ zones where little can be heard.

Many studies optimise the placement of these sensors using a simple detection range calculation based on water column properties of the surrounding area. This is not valid for regions with highly varying bathymetry and strong sound speed gradients where the propagation is non-uniform across the domain. To predict the sonar performance, acoustic propagation simulations are required to accurately determine how the sound energy travels through the region of interest. Studies have previously done this and extended it to Nx2D simulations to create a 3D placement algorithm but still assume their detection range is only as far until the first time the detection probability falls below a threshold [6]. This is also an assumption that should be challenged, as having gone through the computational effort to calculate the propagation loss a more holistic approach can be used looking at the entire domain.

If one were to consider the efficacy of choosing N sensors from a set $\mathcal{S}$ with $|\mathcal{S}| = R$, generally this would involve evaluating R choose N selections of sensors to decide which selection was the best. This set of sensor selections to evaluate grows with $\mathcal{O}(R^N)$ for $N \ll R$, which quickly becomes intractable. To combat this, we leverage the field of submodular optimisation, which gives conditions under which we can perform a greedy selection of sensors in batches of K and still get within a known bound of the optimal solution.




Figure 1: (a) Map of the region of interest off the Victoria coast in Australia. The red line shows the transect used in this study. (b) Example of the propagation loss from a single sensor to a target. The colour bar shows the propagation loss in dB. The red line shows the sound speed profile of the water column.

2. METHODOLOGY

In this paper we select N sensor positions from a set of $|\mathcal{S}| = 2,840$ possible sensor locations from a region of interest off the coast of Victoria, Australia (Fig. 1a), using submodular opti-

mization methods. We further show how choosing sensors in small groups of K can improve proximity to the optimal solution while substantially reducing computation time.

2.1. SUBMODULAR FUNCTIONS

Given a set $\mathcal{S}$ with $|\mathcal{S}| = R$ with subset $S \subseteq \mathcal{S}$, we can define a set function $f(S) : 2^R \rightarrow \mathbb{R}$ that maps the selection S from the set $\mathcal{S}$ onto a scalar. We further let the empty set be denoted $\emptyset$. Set functions have several properties that we leverage and are presented here in brief.

Definition 1. A set function $f(S) : 2^R \rightarrow \mathbb{R}$ is monotone increasing if, for all subsets $A \subseteq B \subseteq \mathcal{S}$ it holds that $f(A) \leq f(B)$.

Definition 2. A set function $f(S) : 2^R \rightarrow \mathbb{R}$ is submodular if, for all subsets $A \subseteq B \subseteq \mathcal{S}$ and $s \notin B$, then

$$f(A \cup \{s\}) - f(A) \geq f(B \cup \{s\}) - f(B). \quad (1)$$

Definition 2 implies adding an element to a smaller set should result in a larger increase to the set function than adding the same element to a larger set (which contains the smaller set). This is often referred to as the set function exhibiting "diminishing returns".

A canonical example of a monotone increasing submodular function stems from the facility location problem.

Example 1 (Facility Location Problem [3]). Let Ω be a set of clients and $\mathcal{S}$ be a set of facilities. Let $d : \Omega \times \mathcal{S} \rightarrow \mathbb{R}$ be the cost of assigning a given client to a given facility. For each client i and subset of facilities $S \subseteq \mathcal{S}$, define $f_i(S) = \max_{s \in S} d(i, s)$. The value of a non-empty S is

$$f(S) = \sum_{i \in \Omega} f_i(S) = \sum_{i \in \Omega} \max_{s \in S} d(i, s), \quad (2)$$

and $f(\emptyset) = 0$. The objective is to select the subset of facilities S of size $|S| = N$ that maximizes $f(S)$. The set function $f(S)$ is a monotone increasing submodular function.

The maximization of a cardinality N set selection problem is NP-hard. However, for the special case of monotone increasing submodular functions, an efficient K -step greedy selection approach performs within 37% of the optimal set selection S^* .

Proposition 1 ([4]). Consider selecting N elements from the set $\mathcal{S}$ which optimizes the value of a monotone increasing submodular cost function $f(S)$ for some $S \subseteq \mathcal{S}$. If we select these elements $K \leq N$ at a time in a greedy fashion where K divides N evenly, then $f(S)$ will satisfy

$$\frac{f(S) - f(\emptyset)}{f(S^*) - f(\emptyset)} \geq 1 - \left(\frac{q-1}{q} \right)^q \geq 1 - \frac{1}{e} \approx 63\% \text{ as } q \rightarrow \infty, \quad (3)$$

where $q = N/K \in \mathbb{Z}^+$.

In practice, the K -step greedy algorithm often performs much better than these bounds. This result also extends to cases where K does not evenly divide N , but we do not consider them in this paper. A graphical interpretation of these results is shown in Fig. 2a. We can also get an intuitive understanding of the computation required to solve these problems optimally by examining Fig. 2b. As an example, it takes approximately 1,000 times more work to evaluate the $K = 3$ -step greedy selection for a set of size $R = 850$ than it does to evaluate the $K = 2$ -step greedy selection from a set of size $R = 500$.




Figure 2: (a) Theoretical submodularity selection bounds on optimality from Eq. 3. If we want to select a total of N elements in batches of K , this shows us how close the final set will be to the performance of choosing N elements in a single batch of N . (b) Relative work to evaluate R -choose- K element selections. This represents how much relative computation is required to select a single batch of K elements from a set of size R . Each line represents the same amount of work. Lines are drawn on each order of magnitude of increasing work.

2.2. PROBLEM STATEMENT

A set of 2,840 possible sensor locations is generated using a uniform distribution across the region of interest from a depth of 0-2000m. From this, we use the MARLIN digital twin (developed on from work in [1]) to simulate the propagation loss using the PyRAM propagation model from every possible sensor location across the domain. PyRAM is a range-dependent acoustic propagation model that uses the split-step Pade method to calculate the propagation loss from a source across 2D environment [2]. As propagation loss can be assumed to be reciprocal, that is the propagation loss from a sensor to a target is the same as from the target to the sensor, we can calculate the propagation loss from every possible sensor location to every point in the region of interest. We take environmental data from the Copernicus Marine Service, which provides water column properties such as temperature, salinity, and pressure. Bathymetry data is taken from the GEBCO dataset, which provides global bathymetry data [5]. The probability of detection is then calculated using the signal-to-noise ratio (SNR) and a logistic function to map the SNR to a probability of detection.

Let Ω be the set of points of interest in the region and $\mathcal{S}$ be the set of possible location. For each point of interest i and sensor s , let the probability of detection be $d(i, s)$. Selecting a subset of N sensors $S \subseteq \mathcal{S}$ that maximizes the average detectability $f(S) = 1/|\Omega| \sum_{i \in \Omega} \max_{s \in S} d(i, s)$, is an instance of the Facility Location Problem from Example 1. This implies that the average detectability $f(S)$ of a set of sensors is monotone increasing submodular function, so a K -step greedy selection of N sensors will be near optimal according to the bounds in Eq. 3.

3. RESULTS

In this section, we will examine the output of a submodular sensor selection problem to a transect off the Victorian coast. If we choose the one sensor that maximizes the average detectability, we can see the output in Fig. 5a. If we instead want to choose two sensors out

of the 2,840 possible sensor locations (see Fig. 3), we would need to evaluate 2840 choose 2, which is approximately 4×10^6 possible pairs. We can see the output of the best selection in Fig. 5d, with an average detectability of $\mu = 0.69$. We leveraged our submodularity results (Eq. 3) to instead choose two sensors sequentially knowing that in the worst case we will get within 37% of the optimal value (see Fig. 2a). Choosing sensors this way will result in evaluating only 5679 sensor placements which is almost 700x fewer evaluations compared to the optimal selection. Figure 5c shows that our average detectability using this method is $\mu = 0.68$, which is within 1.5% of the true optimal we obtained from a brute-force search seen in Fig. 5d. Figure 5b also shows the theoretical lower bounds for $K < N$ from Eq. 3.




Figure 3: Sensor locations: uniformly distributed between 0-2km depth along the entire domain.

We also evaluated selecting three sensors simultaneously, which is 2,840 choose 3 or approximately 3.8×10^9 possible triples. Using the submodular approach selecting individual sensors sequentially we need to evaluate 8,517 sensor placements, which is more than 447,000x faster. One might be tempted to perform a random sensor selection since computing the optimal value for large selections is difficult. We can see in Fig. 4 that selecting $N = 4$ sensors randomly resulted in an average detection probability of $\mu = 0.74$ compared to the average detection probability $\mu = 0.88$ of the greedy selection $K = 1$. In Fig. 5b, we can see how the average detectability increases with the number of sensors selected N . We also see that we do better as the sensors are chosen in larger groups K , up to the globally optimal value of $N = K$.




Figure 4: (a) Single random selection of $N = 4$ sensors with $\mu = 0.75$. (b) Average detectability of a random selection of $N = 4$ sensors over 50 iterations with $\mu = 0.74$.




Figure 5: Selecting multiple sensors to maximise average detectability. (a) $N = K = 1$, (c) $N = 1, K = 2$, (d) $N = K = 2$, (e) $N = 1, K = 3$, (f) $N = K = 3$, (g) $N = 1, K = 4$, and (h) $N = K = 4$. The compute resources for case (h) was 1357x more than for (g), yet the average detectability of (g) was within 1%. (b) Average detectability for $N = K$ and $N < K$.

4. DISCUSSION

Sequential sensor placement impacts detection performance differently than simultaneous placement, but practical constraints also play a role. For rapid deployment, placing the first sensor at the location with the highest average detectability and then selecting subsequent sensors to maximize the average detectability of the remaining locations enables high detection performance at the start of a deployment. However, this approach does not achieve the same overall detection performance as placing sensors using a simultaneous approach. For long-term deployments, simultaneous placement yields higher detection performance across the region of interest, as shown in Fig. 5b. Placing the sensors simultaneously allows for a more holistic view of the region, as the sensors can be placed in locations that complement each other, rather than just maximizing the average detectability of the remaining locations.

Computation time is a key factor in the sensor placement problem. This study has used 2,840 simulated sensor locations which takes up to a week to compute on a standard desktop computer. This could be significantly reduced by using a high-performance computing cluster to run in parallel, but this is not always available. If one then had to evaluate the performance of every possible sensor placement, this would take an intractable amount of time. Submodular optimization allows us to reduce the number of evaluations required to find a near-optimal solution. Taking the example of selecting two sensors from 2,840 possible locations, we can reduce the number of evaluations from approximately 4×10^6 to 5,679, which is a significant reduction in computation time. To see how quickly the computation gets completely intractable, we can consider selecting 10 sensors from 2,840 possible locations, which results in approximately 9.2×10^{27} comparisons in the optimal case, but only 28,355 comparisons with a batch size of $K = 1$. This reduction in computation time is crucial for practical applications, as it allows for rapid deployment of sensors in the field.

5. CONCLUSION

This study has shown that submodular optimization can be used to select sensor locations in a region of interest to maximize the detection performance of the sensors. Using a submodular approach allows us to reduce the number of evaluations required to find a near-optimal solution, which is crucial for practical applications. In real terms, we can reduce computation time of evaluating sensor locations by a factor of 700x for two sensors and 447,000x for three sensors. This makes the near-optimal deployment of modern multi-sensor deployments feasible in a reasonable time frame. An understanding of the trade-offs between sequential and simultaneous sensor placement is also important, as it can impact the detection performance of the sensors.

6. ACKNOWLEDGEMENTS

This study has been conducted using E.U. Copernicus Marine Service Information (<https://doi.org/10.48670/moi-00016>). Edward Clark is grateful to the UKRI Centre for Doctoral Training in Accountable, Responsible, and Transparent AI (ART-AI) and SEA for supporting his PhD.

REFERENCES

- [1] Edward Clark, Alan Hunter, Olga Isupova, and Marcus Donnelly. Optimising sensor path planning with reinforcement learning and passive sonar modelling. In *7th Underwater Acoustics Conference and Exhibition, UACE 2023*, pages 411–418, 2023.
- [2] Marcus Donnelly. Python adaptation of the range-dependent acoustic model (ram)., June 2022.
- [3] Andreas Krause and Daniel Golovin. *Submodular Function Maximization*, pages 71–104. Cambridge University Press, 2 2014.
- [4] G. L. Nemhauser and L. A. Wolsey. An analysis of approximations for maximizing submodular set functions. *Mathematical Programming*, pages 265–294, 1978.
- [5] Pauline Weatherall, Sheila Caceres Ferreras, Sara DA Cardigos, Natalie Cornish, Sam R Davidson, Boris Dorschel, Hayley Drennon, Vicki Ferrini, Hugh A Harper, Tea Isler, et al. The gebco_2024 grid-a continuous terrain model of the global oceans and land. 2024.
- [6] Xiaohan Zhu, Ye Wang, Zeyu Fang, Lei Cheng, and Jianlong Li. Strategic deployment in the deep: Principled underwater sensor placement optimization with three-dimensional acoustic map. *The Journal of the Acoustical Society of America*, 156(4):2668–2685, October 2024.